REMARKS

By this amendment, claims 1-16 are pending, in which no claim is currently amended, newly presented, canceled, or withdrawn.

The Office Action mailed March 16, 2004 rejected claims 1-6, 8, and 10 under 35 U.S.C. § 103(a) as obvious over *Rastogi et al.* (U.S. 6,205,449) in view of *Cooper et al.* (U.S. 6,079,000), claims 7 and 9-15 under 35 U.S.C. § 103(a) as obvious over *Rastogi et al.* in view of *Cooper et al.* and further in view of *Hapner et al.* (U.S. 5,940,827), claim 14 under 35 U.S.C. § 103(a) as obvious over *Rastogi et al.* in view of *Hapner et al.,* and claim 16 under 35 U.S.C. § 103(a) as obvious over *Rastogi et al.* in view of *Cooper et al.* and further in view of *Nilsen et al.* (U.S. 5,668,986).[1]

The rejection of claims 1-10 over *Rastogi et al.* and *Cooper et al.* (with or without *Hapner et al.*) is respectfully traversed because the references do not teach or otherwise suggest the limitations of the claims. For example, independent claim 1 recites: "synchronizing a transaction performed on the primary database system based on a number of transactions in the buffer and **a predetermined number of transactions**."

Synchronizing a transaction based on a "predetermined number of transactions" advantageously addresses difficulties in database systems. For example, the Background section of the present application explains that "it is difficult to characterize the amount of data lost in terms that database owners can best understand. The maximum exposure for loss of data in this approach is usually described in terms of the size of the redo logs, but this information is not helpful for database owners, who would rather want to know how many orders were lost" (¶ 7).

[1] Notwithstanding the Office Action's inclusion of an apparent response to arguments traversing the previous rejection of claims 10, 13, and 15, the Office Action of March 16, 2004 did not affirmatively set forth a rejection of claims 10, 13, and 15—or any other claim—under 35 U.S.C. § 112, ¶ 2. Therefore, the previous rejection of claims 10, 13, and 15 under 35 U.S.C. § 112, ¶ 2 is considered withdrawn.

In recognition of this problem, the method recited in claim 1 sets forth the use of "a predetermined number of transactions" in synchronizing a transaction. "Because the pre-determined bound is specified in terms of the number of transactions, the database operator can set a meaningful tradeoff between performance and data availability that is appropriate for the particular needs of the database operator's installation" (Summary, ¶ 11).

The Office Action correctly acknowledges that "Rastogi does not explicitly teach a predetermined number of transactions" (p. 3). Indeed, the portion of *Rastogi et al.* cited for transactions, col. 8:3-8, only discusses "transactions" in the model described in the reference. There is no mention or suggestion of "synchronizing a transaction," much less "synchronizing a transaction" based on "a predetermined number of transactions."

Cooper et al. too fails to show this feature. In fact, *Cooper et al.* exhibits many of the same difficulties described in the background and addressed by the invention recited in claim 1. For example, *Cooper et al.* recommends managing the "optimal transfer efficiency" between the audit host memory **342** and the XPC cache area **350** based on a "predetermined size of audit host memory **342**," which is given in terms of a "predetermined number of address locations within audit host memory **342**" (col. 12:39-40). Referring now to FIG. 9 of *Cooper et al.*, transactions **344**, **354**, **362**, and **370** clearly have different sizes in terms of the number of audit host memory **342** locations. For example, transaction **354** is about half the size of transaction **344** and about a third of the size of transaction **362**. When transaction sizes are so variable as in *Cooper et al.*, pre-specifying buffers in terms of number of memory locations or bytes does not meaningfully specify "a predetermined number of transactions" as set forth in independent claim 1. In fact, this variability in transaction size is what dooms *Cooper et al.*'s approach to exhibit the problems addressed by the invention of independent claim 1.

The portion of *Cooper et al.* cited in the Office Action, col. 12:30-43, does not support the rejection. Although *Cooper et al.* speculates that "the optimal transfer characteristics of the physical audit trail **378** may determine when a sufficient number of transactions have been accumulated in audit host memory **342**" (col. 12:33-36), *Cooper et al.* then defines what it means a "sufficient number of transactions"—not in terms of a "predetermined number of transactions" as recited in claim 1—but explicitly in terms of the size of audit host memory **342**: "Thus, the optimal transfer efficiency may correspond to the a predetermined size of audit host memory **342**" (col. 12:36-38). Due to the variability in transaction sizes evident in FIG. 9, *Cooper et al.*'s criterion based on a predetermined memory size can only crudely and inaccurately correspond to the actual number of transactions in the audit host memory **342**. However, *Cooper et al.* is not concerned with providing a meaningful bound to database administrators but focused on transferring data between memories as efficiently as possible. Thus, *Cooper et al.* actually **teaches against** using a "predetermined number of transactions," since transferring one predetermined number of tiny transactions is less efficient than transferring the same predetermined number of large transactions. If a proposed modification would render the reference being modified unsatisfactory for its intended purpose, then there is no suggestion or motivation to make the proposed modification. *In re Gordon*, 733 F.2d 900, 221 USPQ 1125 (Fed. Cir. 1984).

As for col. 14:41-43, also cited in the Office Action, *Cooper et al.* makes a similar recommendation for the size of the XPC cache area **350** in terms of "a predetermined number of address locations within XPC cache area **350**" (col. 14:39-41) and is therefore similarly deficient for the same reasons as with the size of audit host memory **342**. Furthermore, *Cooper et al.* states that the "synchronous audit data request results in a portion of the audit host memory **342** being written to a corresponding portion of non-volatile memory storage such as XPC cache **350**"

(col. 11:50-53), thus the relevance of the size of the XPC cache area **350** to "synchronizing a transaction" as recited in independent claim 1 is not immediately apparent.

Concerning dependent claims 7 and 9 as well as claims 11-16, the use of the non-analogous *Hapner et al.* does not support the rejection by curing the deficiencies of *Rastogi et al.* and *Cooper et al.* or by disclosing the additional features recited in claims 7, 9, and 11-16. *Hapner et al.* relates to maintaining a database cache in conjunction with a persistent database portion and uses a "transaction counter" (col. 3:44) for keeping track of how many open transactions there are in the database cache **140**. *Hapner et al.* lacks any disclosure of using such a "transaction counter" for any set of "transactions to be sent to a standby database system." Since neither *Rastogi et al.* nor *Cooper et al.* operate with a predetermined number of transactions to be sent to a standby database system, there is no motivation to modify *Rastogi et al.* and/or *Cooper et al.* to count something none of the references seem to care about.

Furthermore, the only comparison of *Hapner et al.*'s transaction counter is with zero (0) in FIG. 10, step 469, and FIG. 11, steps 512 and 536. However, claim 11 explicitly recites "**compare** the counter and **the predetermined bound**" and claim 12 recites "**comparing** a counter indicating a number of the transactions in a queue of transactions to be sent to a standby database system and **a predetermined bound of transactions**." Whatever *Cooper et al.* may be thought to disclose about the optimal transfer size of audit host memory **342**, the optimal transfer size certainly cannot be zero! Thus, even if the Examiner might be correct in responding that "Hapner does apply a transaction counter in a replicating data" (p. 11), the claims are more specific than that and the use of a counter in the specific context of claims 7, 9, and 11-16 is not disclosed in *Hapner et al.*

Nilsen et al., applied only against claim 16, does not furnish a disclosure of "synchronizing a transaction performed on the primary database system based on a number of transactions in the buffer and the corresponding bound" which is missing in *Rastogi* and *Cooper et al.* Thus, claim 16 too is patentable over *Rostogi et al.*, *Cooper et al.*, and *Nilsen et al.*

Therefore, the present application, as amended, overcomes the objections and rejections of record and is in condition for allowance. Favorable consideration is respectfully requested. If any unresolved issues remain, it is respectfully requested that the Examiner telephone the undersigned attorney at 703-425-6499 so that such issues may be resolved as expeditiously as possible.

Respectfully Submitted,

DITTHAVONG & CARLSON, P.C.

5/17/2004
Date

Stephen C. Carlson
Reg. No. 39,929

Attorney for Applicant(s)

10507 Braddock Rd
Suite A
Fairfax, VA 22032
Tel. 703-425-6499
Fax. 703-425-8518